

Collaborative zkSNARKs

Zero-knowledge proofs for *distributed* secrets

Alex Ozdemir, Dan Boneh

To appear in USENIX Security'22

Provable properties about secrets

Authentication



zkSNARKs

I **prove** that I **know** my
password/secret key.

Client(secret)


Server



Provable properties about *distributed* secrets


Money Laundering

Several banks prove that “Eve” has Colonial Pipeline’s ransom.





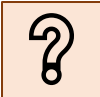
Healthcare Statistics

Several hospitals prove that procedure prices are “fair”.



This talk

Applications:

- Authentication  zkSNARKs
- Money laundering 
- Healthcare statistics 

Outline:

1. Why zkSNARKs are insufficient
2. New tool: ***collaborative* zkSNARK**
3. Building collaborative zkSNARKs
4. Surprising efficiency

Background

zkSNARKs

Witness relations

$$\underline{P(x, w)}$$

$$\underline{V(x)^{00}}$$

$$\xrightarrow{w} (x, w) \in R$$

$$\downarrow \\ \{0, 1\}$$

$$\{ \exists? w. (x, w) \in R \}$$

- w may be large
- not private

zkSNARKs

$pp \leftarrow \text{Setup}()$

$P(x, w)$

$\pi \leftarrow \text{Prove}(pp, x, w)$

π

$\text{Verify}(pp, x, \pi)$


\downarrow
 $\{0, 1\}$

$V(x)$

$\{ \exists? w. (x, w) \in R \}$

- Sound: π proves w exists
- (zk) zero knowledge: hides w
- (S) succinct: short π , fast Verify
- (N) non-interactive
- (AR) argument: *computationally* sound
- (K) knowledge: P knows w

Authentication with a zkSNARK



Existing zkSNARKs


- From pairings & elliptic curves
 - Groth16
 - Marlin (KZG)
 - Plonk (KZG)
 - ...
- From hashing & codes
 - Fractal
 - ...
- ...

*For distributed secret data,
who plays the prover?*

Collaborative zkSNARKs

Definitions

Collaborative zkSNARKs




Syntax:

- $\text{Setup}() \rightarrow pp$
- $\text{Prove}(pp, x, P_1(w_1), \dots, P_N(w_N)) \rightarrow \pi$
- $\text{Verify}(pp, x, \pi) \rightarrow \{0, 1\}$

t-zero-knowledge

Any adversary controlling $\leq t$ provers learns nothing but whether

$$(x, \vec{w}) \in R$$




- Formally

- Adversary corrupts $\leq t$ provers
- ZK simulator is given
 - the corrupt witnesses
 - x
 - $b = (x, \vec{w}) \in R$
- We use the *random oracle model*

Knowledge soundness

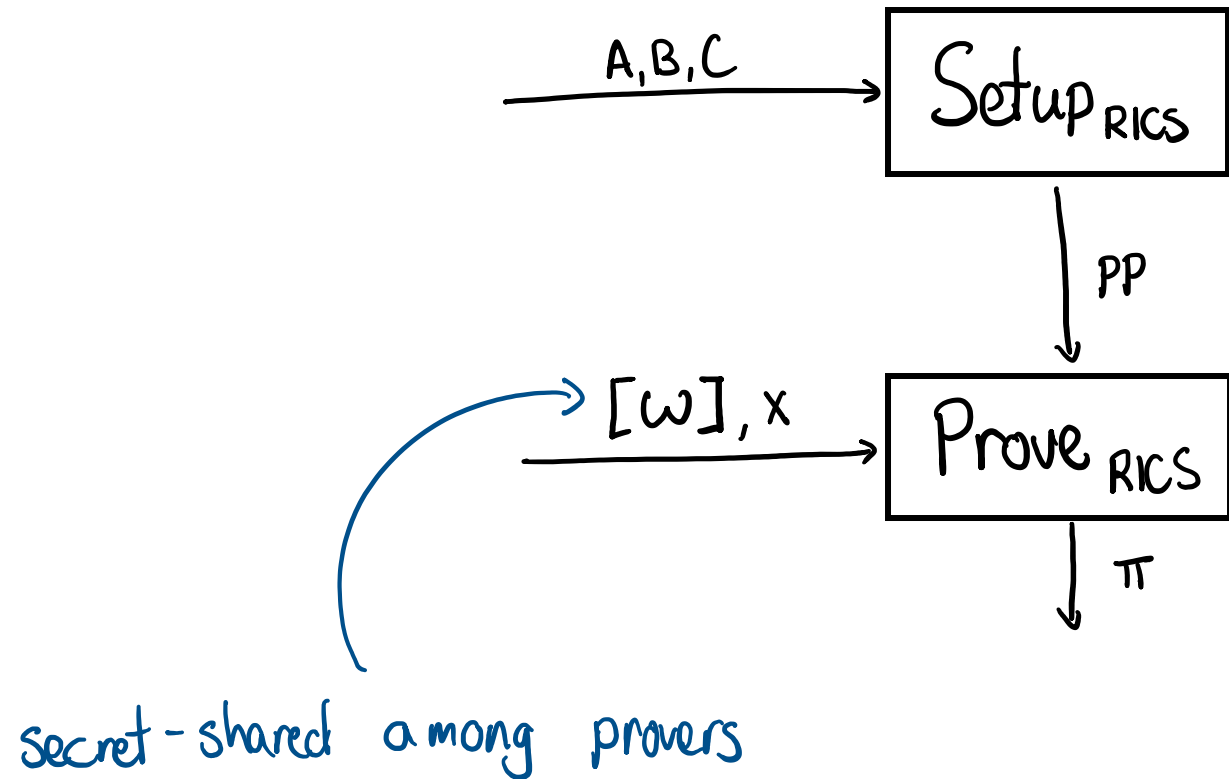
- *could* mean P_1 knows w_1, \dots, P_N knows w_N
- actually means P_1, \dots, P_N collectively know w_1, \dots, w_N
 - “distributed knowledge” [Halpern, Moses '90]
- Random oracle
 - extractor programs RO



Our focus: secret-shared R1CS witnesses

R1CS:

- Class of relations
- Generalize arithmetic circuits
- Definition:
 - $R: A, B, C \in \mathbb{F}^{n \times m}$
 - $x \in \mathbb{F}^k, w \in \mathbb{F}^{m-k}$
 - Satisfied when
 - $a \leftarrow w \parallel x$
 - $Aa \circ Ba = Ca$



Designing co-zkSNARKs

Overview of constructions


Approach: MPC the Prover

GenericMPC(zkSNARK.Prove, $[w]$) $\rightarrow \pi$

$\underbrace{\hspace{1.5cm}}_{1000\times \text{ slower}} \underbrace{\hspace{1.5cm}}_{1000\times \text{ slower}}$

$1,000,000\times$ slowdown ?!

avoid. instead achieve $1000\times - 2000\times$ slowdown



Potential Bottlenecks

Single-prover bottlenecks:

- Elliptic curve operations
- Fourier transforms

? This talk: a good solution

✓ MPC-efficient

MPC bottlenecks:

- Polynomial divisions
- Partial products
- Merkle tree evaluations


✓ MPC-efficient (for SNARK provers)

✓ special protocol

? This talk: an okay solution

MPC Crash-Course

Computation: arithmetic circuit over a finite field



1. Secret-share wire values among N parties

e.g. $x = \underbrace{x^{(1)} + x^{(2)} + \dots + x^{(N)}}_{[x]}$

2. Secure protocols for +, * on shares
3. Evaluate circuit, inputs to outputs

MPC-friendly elliptic curve arithmetic

Option 1: Share (x, y) coordinates

Elliptic curve addition

$$g_1 = (x_1, y_1) \in E(\mathbb{F})$$

$$g_2 = (x_2, y_2)$$

$$g_1 \boxplus g_2 ?$$

$$[g_1] \boxplus [g_2] =$$

$$([x_1], [y_1]) \boxplus ([x_2], [y_2]) \rightarrow$$

$$x = \frac{y_1 y_2 + x_2 y_1}{1 + d x_1 x_2}$$

$$y = \frac{y_1^2 - d x_1^2 y_2}{1 - d x_1 x_2 y_1}$$

Nasty Formulas

Expensive

- multiple multiplications
- communication

Option 2: Elliptic curve sharing


$$[g] = g^{(1)} \boxplus g^{(2)} \boxplus \dots \boxplus g^{(N)} \rightarrow$$

add
share-wise

Efficient

- no communication



Merkle Trees



non-linear and expensive



Idea (Riad S. Wahby): reduce
with proof composition



proof size
(& verifier work)
 $\times N$

Implementation

Implementation Goals

- Three base zkSNARKs
 - Groth16
 - Marlin/KZG
 - Plonk/KZG
- Two base MPCs
 - GSZ: $t < N/2$ (honest majority)
 - SPDZ: $t < N$ (malicious majority)
- Goals:
 - Compete with existing zkSNARKs
 - (well optimized!)
 - Iterate on MPCs, sub-protocols
 - Don't work too hard

An Opportunity

1. Arkworks has curve-generic provers:


```
fn prove<E: PairingEngine>(..) {  
    ...  
}
```

2. Curve interfaces define +, *, ...

```
trait PairingEngine {  
    type ScalarField;  
    type Curve;  
    fn field_add(...) ...;  
    fn field_mul(...) ...;  
    fn curve_add(...) ...;  
    ...  
}
```

Implementation Strategy

1. Implement MPCs for shared field and curve operations
 1. SDPZ
 2. GSZ
2. Wrap MPCs & implement arkworks interfaces
3. Instantiate zkSNARK prover
 - Mis-appropriates zkSNARK prover as a co-zkSNARK prover!



Performance

Experimental Setup

Measure:

- Wall-clock proving time

Vary:


- N : number of provers
- n : R1CS size (# constraints)
- c : link capacity
- Base: Groth16/Marlin/Plonk
- t : security threshold
 - $< N/2$ (honest majority, GSZ)
 - $< N$ (malicious majority, SPDZ)

Simplifications:

- No intra-prover parallelism
- Skip MPC preprocessing
 - Small for our computation

Experiment 1: Good network, few parties

Fix a 3Gb/s link, vary # rank-1 constraints

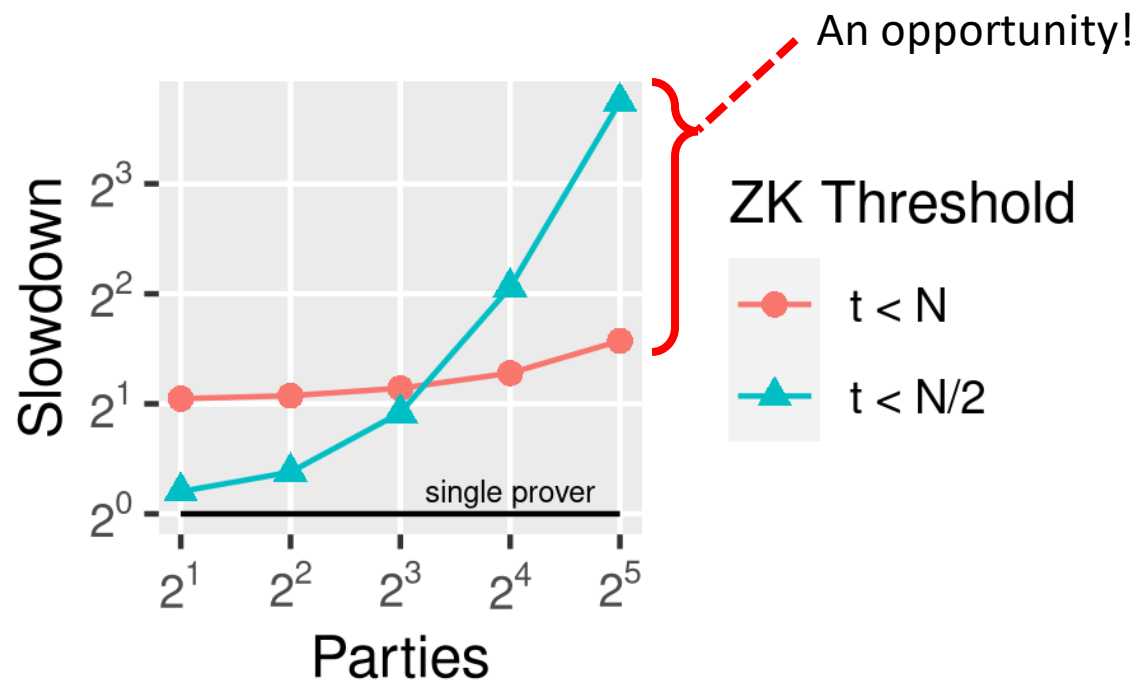


$t < N/2 \rightarrow$ no slowdown

$t < N \rightarrow$ 2x slowdown

Experiment 2: Many provers

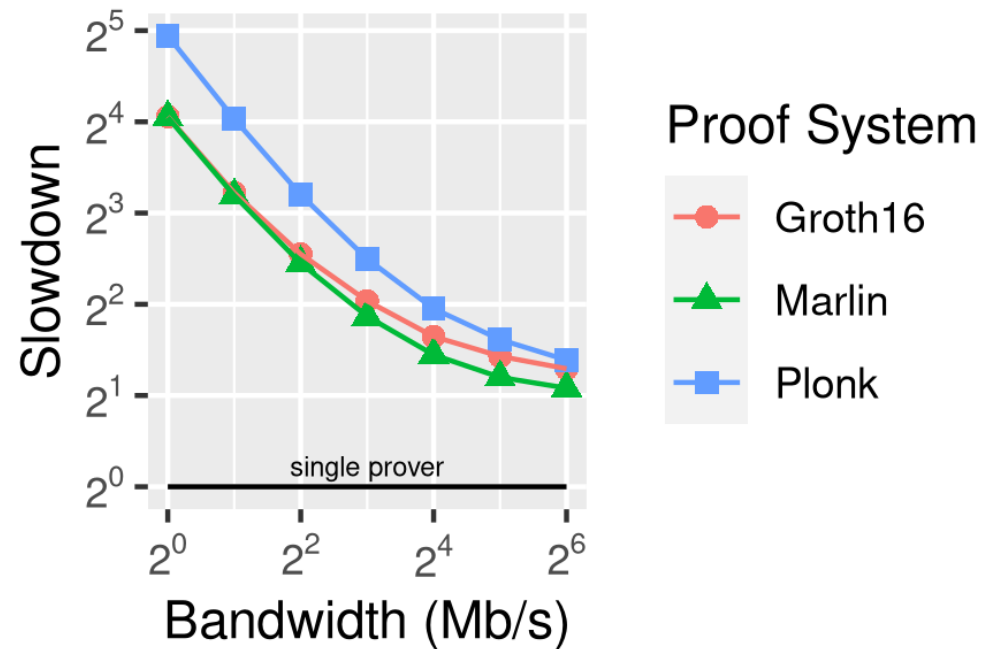
Fix 1024 constraints, 3Gb/s link, Groth16, vary # of provers



Slowdown grows with N ; better for SPDZ

Experiment 3: Low-capacity link

Fix 1024 constraints, 2 provers, malicious majority (SPDZ)



Slowdowns grow, but *far* better than 1000x

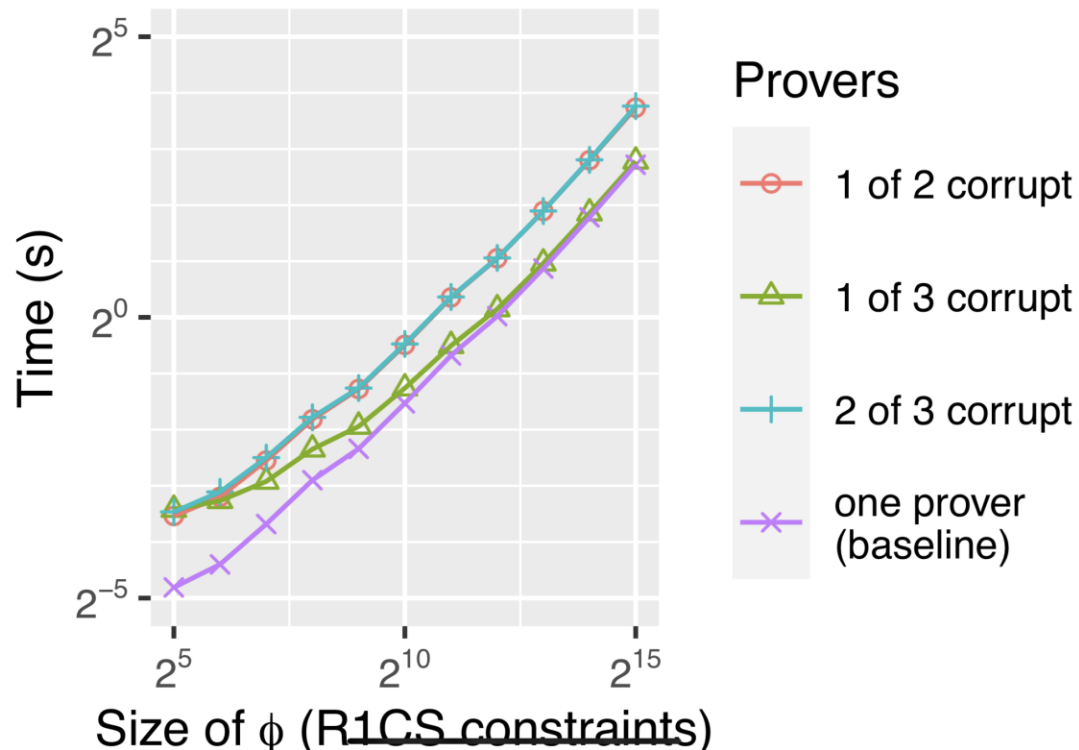
Discussion, Future Work

- Bandwidth is the bottleneck for many provers, low link capacity
 - **Bad news:** 2-prover co-zkSNARK (additive sharing) $\rightarrow \Omega(n)$ communication
 - From randomized 2-party communication complexity of DISJOINT
 - Conjecture: $\Omega(\lambda n)$ (needed: generalize DISJOINT from $\{0,1\}$ to \mathbb{F})
- Exploit intra-prover parallelism
- A nicer post-quantum co-zkSNARK with $o(N)$ proof-size?

Collaborative zkSNARKs

Alex Ozdemir, Dan Boneh

Groth16 co-zkSNARK proving time



Conclusions:

1. **Collaborative zkSNARKs** support **distributed secrets**
 - Multiple users, hospitals, banks, ...
2. Very efficient
 - $(N/2)$ -ZK \rightarrow **no slowdown**
 - $(N - 1)$ -ZK \rightarrow **2x slowdown**
3. Far better than MPC for typical computations \rightarrow **$\sim 1000x$ slowdown**

Code

Component	Lines (Rust)
Network Library	~700
Arkworks adapters	~2000
MPC protocols	~3000
Plonk	~1200